

# THE DEEP SPACE 1 EXTENDED MISSION

Marc D. Rayman and Philip Varghese  
Jet Propulsion Laboratory  
California Institute of Technology  
4800 Oak Grove Dr.  
Pasadena, CA 91109 USA

The primary mission of Deep Space 1 (DS1), the first flight of the New Millennium program, completed successfully in September 1999, having exceeded its objectives of testing new, high-risk technologies important for future space and Earth science missions. DS1 is now in its extended mission, with plans to take advantage of the advanced technologies, including solar electric propulsion, to conduct an encounter with Comet 19P/Borrelly in September 2001. During the extended mission, the spacecraft's commercial star tracker failed; this critical loss prevented the spacecraft from achieving three-axis attitude control or knowledge. A two-phased approach to recovering from the loss of the critical star tracker was undertaken. The first involved devising a new method of pointing the high-gain antenna to Earth using the radio signal received at the Deep Space Network as an indicator of spacecraft attitude. The second was the development of new flight software that allowed the spacecraft to return to three-axis operation without substantial ground assistance. The principal new feature of the software is the use of the science camera as an attitude sensor; to accomplish this, the new system builds upon some of the software used by the autonomous optical navigation system that was tested and used operationally during the primary mission. The differences between the science camera and the star tracker have important implications not only for the design of the new software but also for the methods of operating the spacecraft and conducting the mission. The ambitious recovery was fully successful, and the mission is back on track.

## INTRODUCTION

Deep Space 1 (DS1) was launched on October 24, 1998 as the first mission in NASA's New Millennium program (NMP). NMP is designed to accelerate the realization of ambitious missions by developing and validating some of the high-risk, high-benefit technologies they need. NMP conducts deep space and Earth orbiting missions focused on the validation of these technologies.<sup>1</sup> Thus, the 11-month primary mission of DS1 was devoted to the testing and evaluation of 12 technologies selected on the bases of their importance to future space and Earth science programs, the significant advancement they offer over current state-of-the-art, the high risk they present to the first user, and the need for in-flight testing to reduce that risk.

In addition to its technical objectives, DS1 was intended to probe the limits of rapid devel-

opment for deep-space missions. The initial study of DS1 was undertaken only 39 months before launch, an unprecedentedly short time for a NASA deep-space mission in the modern era. At the time the preliminary concept study was initiated, the only definition of the project was that it would validate solar electric propulsion and other unidentified technologies in deep space and that launch would occur sometime in 1998. The level-1 requirements and goals were formulated 26 months prior to launch.

The 12 advanced technologies tested during the primary mission are:

- solar electric propulsion
- solar concentrator arrays
- miniature integrated ion and electron spectrometer
- miniature integrated camera and imaging spectrometer
- autonomous optical navigation
- small deep-space transponder
- K<sub>a</sub>-band solid state power amplifier
- beacon monitor operations
- autonomous remote agent
- low-power electronics
- power actuation and switching module
- multifunctional structure

---

Copyright 2000 by the American Institute of Aeronautics and Astronautics, Inc. The U.S. Government has a royalty-free license to exercise all rights under the copyright claimed herein for Governmental purposes. All other rights are reserved by the copyright owner.

The technologies and their performances during the primary mission are described by Rayman et al.<sup>2</sup> and references therein. All but the last four technologies listed above play some role in the extended mission. Some of them, including the solar arrays, transponder, and the real-time part of the autonomous navigation system (which provides positions of the Sun, Earth, and other bodies to the attitude control system), are critical to basic spacecraft functioning.

All of the technology testing was completed by July 1999. Following the end of the testing, DS1 conducted a bonus encounter with asteroid 9969 Braille on July 29, 1999. The primary mission completed on September 18, 1999, having met or exceeded all of the mission success criteria. The total cost, including launch but excluding the development cost of some of the technologies, from the formulation of the project through the end of the mission was \$149.7 M (in real-year dollars). Further background on the Deep Space 1 project, the mission, and the spacecraft has been presented elsewhere.<sup>1-3</sup>

### MISSION PROFILE

When DS1 was launched, the plan for the primary mission incorporated the encounter with Braille, although it was not a required part of the mission. The mission design maintained an option for an extended mission encounter with Comet 19P/Borrelly in September 2001. After launch the mission progressed so well, with the critical technologies exhibiting excellent performance, that the proposal for the extended mission was augmented to include an encounter with Comet 107P/Wilson-Harrington. The extension to the mission was approved by NASA in August, 1999. That extended mission was described by Rayman et al.<sup>2</sup>

Thrusting with the ion propulsion system (IPS) during the primary mission was designed to allow extensive testing of the technology. It also placed the spacecraft on the trajectory to Braille and the extended mission targets. Although the request for the extended mission had not yet been granted, thrusting resumed only about 36 hours after the closest approach to Braille so that if the proposal were approved, the spacecraft would remain on course for the comets.

A typical week of IPS thrusting began with the autonomous navigation system (AutoNav) commanding the spacecraft to turn to point the ion engine in the direction required for optimal thrusting. AutoNav then started the IPS and throughout the week updated the thrust direction and throttle level. After about 150 hours, the IPS was turned off and AutoNav commanded ACS and the miniature integrated camera/spectrometer (MICAS) to collect visible images of selected asteroids and background stars for its use in on-board orbit determination. The collection of such images could last for up to 4 hours, at the end of which AutoNav would point the unarticulated high-gain antenna (HGA) to Earth for the weekly track by the Deep Space Network (DSN). At the end of the DSN session, AutoNav would take the spacecraft back to the thrust attitude. The workload for the operations team was significantly less than it would have been without AutoNav.

Optimal trajectories that use solar electric propulsion have periods in which coasting is better than thrusting. Because the encounter targets were not required for the DS1 mission, the timing of thrust and coast periods was a critical part of the selection of the mission design. Only targets that allowed coast periods at times that matched the needs of the extensive technology testing program during the primary mission were considered. Some coasting periods in the thrust profile were optimal and others were inserted to allow special spacecraft activities or to provide buffer against possible periods of unexpected loss of thrust.

Following nearly three months of thrusting after the Braille encounter, a coast period began on October 20, 1999. At that point, the IPS had consumed 21.6 kg of the 81.5 kg of xenon on board at launch, yielded 1320 m/s, and completed 3571 hours of operation.

During the coast period, planned to last until the middle of December 1999, extensive calibrations of all channels in MICAS were conducted. In addition, MICAS acquired 48 infrared spectra of Mars, covering nearly two full rotations of the planet, from a range of 55 million km. With three spectra collected every three hours for 48 hours, the data allow resolution of 45° in longitude. These data are considered to be the highest quality spectra of Mars in the range of 1.3  $\mu$ m to 1.9  $\mu$ m.

Spectral features have been detected which may indicate the presence of previously unrecognized surface minerals.<sup>4</sup>

### STELLAR REFERENCE UNIT FAILURE

On November 11, 1999, after all the Mars spectra were acquired but before the next scheduled DSN pass, at which the data would be returned, the spacecraft's stellar reference unit (SRU) stopped reporting attitude data to the spacecraft computer. This commercial unit recognized star patterns and produced a quaternion and thus was able to provide the complete three-axis attitude. The SRU was one of three attitude sensors; the spacecraft also carries one laser gyro for each axis and a Sun sensor assembly (SSA), with 128° full-angle field of view, that was used principally for safe modes. Attitude is controlled using a hydrazine-based reaction control system; when the IPS is thrusting, ACS controls two axes by moving the ion thruster through a range of  $\pm 5^\circ$ . Thrust vector control (TVC) using the IPS substantially reduces hydrazine expenditure.

The SRU had exhibited intermittent problems since shortly after launch. Diagnostic activities on board and laboratory tests conducted jointly with the vendor and JPL had not yielded an explanation for the occasional interruptions in its reporting of attitude. The longest outage was 28 minutes.

When the SRU exhibited problems on November 11, the spacecraft's fault protection system power cycled it two times before finally declaring a celestial inertial reference loss (CIRL). CIRL leads to the spacecraft entering a safe mode known as Sun standby SSA. In this mode, the SRU and other devices are power cycled, non-essential devices are turned off, and ACS uses the SSA and gyros to point the spacecraft +x axis at the Sun and rotate around the Sun-spacecraft line at 1 revolution/hour. The center of the spacecraft's high gain antenna and the center of the SSA's field of view are along the +x axis. ACS also rotates the solar arrays so that they are normal to x.

The spacecraft has three low-gain antennas (LGAs): one each aligned with +x, +z, and -z. The HGA and LGAs all work at X-band. (There is also a +x K<sub>a</sub>-band antenna that has been used principally for technology validations

and for DSN testing but can also return telemetry to the DSN stations equipped for K<sub>a</sub>-band reception.) The rotation in Sun standby SSA is a remnant from a very early mission phase in which the Sun-probe-Earth (SPE) angle was too large to return telemetry through the +x LGA. For the remainder of the mission, the SPE will remain less than 45°, so the +x LGA will always be used in this safe state; at the time of the SRU failure, the SPE angle was 38°. At a geocentric range of 1.6 astronomical units (AU), the 34-m stations of the DSN were unable to close the link for downlink, so no telemetry could be returned, although they were capable of commanding the spacecraft. The 70-m stations could support 150 bits/s downlink, but they did not have X-band uplink. (Since then, the 70-m station at Goldstone Deep Space Communications Complex has been augmented with X-band uplink, and upgrades for the 70-m stations at the Canberra and Madrid Deep Space Communications Complexes are scheduled to be completed in November 2000 and October 2001 respectively.)

The initial analysis of the SRU failure was severely limited by the low downlink rate and the limited DSN coverage that had been scheduled. A large volume of engineering telemetry to provide the complete context for the SRU's anomalous behavior needed to be returned. Before the return of all the data, several attempts to revive it were conducted, all without success. By the end of November, evidence was growing that the SRU could not be restored.

During development the SRU was considered a critical spacecraft device. Funding was not adequate however to provide for redundancy. As a result, the loss of the SRU was considered to be a mission-ending failure. Nevertheless, the project undertook an extremely rapid and extensive recovery effort in two phases.

### PHASE 1 RECOVERY

It was clear that to conduct a thorough diagnosis of the SRU, to return the large volume of Mars data, and to conduct any further meaningful activities with the spacecraft, it would be necessary to use the HGA. Thus, the first phase of a recovery was initiated, with the objective of pointing the HGA to Earth. Based

on rapid testing in the DS1 testbed at JPL and several tests conducted with the spacecraft, an experimental procedure was developed and executed successfully on January 14, 2000.

With only gyros and the SSA, ACS has knowledge of the Sun location but not of other celestial bodies. To point the HGA to Earth, the first step was to command the spacecraft to offset the Sun from the center of the SSA by the SPE angle. On January 14, that angle was  $34.3^\circ$ . Although the direction of offset could be specified in spacecraft body coordinates, the relationship of that direction to inertial space was unknown. Once the offset was achieved, the rotation rate around the Sun-spacecraft line (now  $34.3^\circ$  from the  $+x$  axis) was commanded to 1 revolution/45 minutes. The spacecraft transmitted an unmodulated carrier through its HGA as it coned around the Sun.

In Sun standby SSA, the solar arrays are caged so that the plane of the panels is normal to  $+x$ . Because the arrays have cylindrical concentrator lenses, they are very sensitive to pointing in one axis; thus, extra commands were included to articulate the arrays by the SPE angle so that they would remain orthogonal to the Sun-spacecraft line.

The DSN observed the X-band signal as the spacecraft eventually swept past Earth, revealing the unknown phase of the rotation. The peak of the received signal corresponded to the HGA being Earth-pointed. Because of uncertainty in gyro bias values, two peaks were used to refine the frequency of the rotation. There was also uncertainty in the time it would take ACS, still operating in its Sun standby SSA mode, to achieve steady-state rotation, so the first peak was observed but not used for measuring the frequency.

Once the phase and frequency of the coning were known, the time of the next peak was predicted. A special short uplink frequency sweep had been developed with the DSN, and, accounting for the one-way light time (13 minutes 55 seconds on January 14), the beamwidth of the HGA ( $4^\circ$  from the center to the 3 dB point), and the rotation rate, the time of the beginning of the sweep was computed. The sweep would begin at a time that would make it arrive at the spacecraft as the leading edge of the HGA moved Earth into its beam. The uplink

sweep completed, thus bringing the spacecraft receiver into frequency lock, in time to allow one command to be uplinked before the rotation would take the trailing edge of the HGA out of view of Earth.

The single command that was transmitted activated a sequence that had been uplinked earlier through the LGA. The sequence commanded the spacecraft to stop coning around the Sun. Analyses and tests had enabled predictions for the length of the delay caused by the sluggish response of the ACS in this mode as well as the expected interval past the time of  $+x$  on Earth that the command would be received (given the uplink sweep time, command radiation time, and other delays in the system). Thus, the sequence included commands for the spacecraft to rotate back far enough to account for these effects. When the next peak was observed, the signature of the carrier power at the DSN clearly revealed the spacecraft continuing past it and eventually backing up, ending up with the HGA within 2 dB of the predicted value for optimal Earth-pointing.

Once the HGA was pointed to Earth, gyro drift would gradually move it away. ACS would keep the Sun at the desired angle from the center of the SSA, so the effect of the gyro drift was to continue move the spacecraft around the Sun-spacecraft line. Ten sequences, each commanding the spacecraft to rotate a fixed angle from  $-10^\circ$  to  $+10^\circ$  in  $2^\circ$  increments, were stored on board. As the observed carrier power at the DSN diminished to the point that a correction was deemed necessary, one of the 10 sequences would be activated by real-time command, based on how much the signal had decreased and whether it was seen to have passed through the peak before decreasing (thus indicating the sign of the correction that was needed).

Real-time commands were used to select the uplink and downlink rates, based on how accurate the HGA pointing was. At the end of a DSN session, the spacecraft was commanded to point  $+x$  to the Sun again. If it were left at the SPE offset angle, gyro drift during the gap in DSN coverage (which was typically 1 week) would have been sufficient to make the spacecraft attitude and, therefore, the LGA direction, unpredictable. With  $+x$  pointed at the

Sun, the uncertainty in phase around the Sun-line was irrelevant for LGA communications and the initiation of subsequent repointing activities.

Once the HGA was Earth-pointed, the first priority was return of the Mars spectra. With DS1's future being in serious doubt, given the inability to revive the SRU, it was considered most important to return the science data that had already been acquired.

The capability to point the HGA to Earth allowed a more complete investigation into the state of the SRU to begin. Extensive diagnostic activities were conducted that simply would have been far too data intensive through the LGA. As JPL worked with the contractor that provided the unit, it became evident that its failure was permanent. Despite a significant effort, however, the failure mechanism could not be established.

The pointing procedure proved extremely successful and productive and was used frequently from January through June. Once the HGA was Earth-pointed, attitude corrections needed to be transmitted only about once every 2 hours. Nevertheless, the procedure did consume valuable DSN time during the coning and each activity required diligence and planning (to account for changing SPE angle and geocentric range) that was not negligible for the very small DS1 operations team. In addition, although the spacecraft could be controlled to point +x to Earth, this technique did not permit a practical way to achieve any other attitude.

### PHASE 2 RECOVERY

The permanent loss of the SRU meant that if any further worthwhile operations were to be attempted with the spacecraft, a new method of controlling the attitude would be necessary. Even returning to technology validation activities would be impractical without a change. For example, conducting further tests with the IPS would produce a small torque that ACS would not be able to counter with the ion thruster gimbal in the absence of attitude knowledge. By the end of January, about 16 kg of hydrazine (from an initial load of 31 kg) remained on board; this was insufficient to control the attitude for more than a few days of IPS thrusting. Thrusting without using thrust vector control would be extremely costly.

The DS1 project elected to attempt a complete recovery of the spacecraft in time to provide an opportunity to conduct a comet encounter. Throughout development and operations a considerable body of work had been devoted to analyzing the trajectory DS1 was planning to follow; no low-thrust trajectory had ever been studied in such detail. An important figure of merit for a low-thrust mission has been determined to be the susceptibility to unexpected missed thrusting. While anomalies short enough to cause significant problems for a conventional chemical propulsion mission (such as missing a major trajectory correction maneuver) are unlikely to have much effect on a low-thrust mission, the mission may still be susceptible to long-periods of missed thrusting. Several techniques were employed to build margin into DS1's trajectory, and it could accommodate periods of well over a month of lost thrust. The extensive phase 2 recovery operations however would exceed the time that the spacecraft could miss IPS thrusting and still reach both targets. As a result, it became necessary to abandon at least one of the comets in the extended mission. The Deep Space 1 Science Team selected the original extended mission target, Comet Borrelly, over Comet Wilson-Harrington. To reach Borrelly in September 2001, IPS thrusting had to resume by late July 2000.

In January 2000, in parallel with detailed investigations into the SRU, several methods for replacing the attitude knowledge that had been provided by the SRU were considered, but the one that was selected relied upon using the visible CCD channel in MICAS to track a star for attitude reference. Differences between MICAS and the SRU made this replacement far from straightforward however. Table 1 shows some of those differences.

The design was complicated by the presence of scattered light in MICAS. The scattered light was studied extensively as part of the mission's technology validation experiments and was well understood from in-flight testing and modeling. In many attitudes it reduced the useful field of view by about 30% from what is shown in Table 1. The combination of scattered light and regions of the MICAS field of view (FOV) with decreased sensitivity limit the

faintest star that can be used for attitude reference to  $m_v \approx 6$ .

Parameter	SRU	MICAS
Field of view	$8.8^\circ \times 8.8^\circ$	$0.69^\circ \times 0.78^\circ$
Limiting stellar magnitude	$m_v = 7.5$	$m_v = 9.5$
Output format	Quaternion	Image file
Output rate to spacecraft computer	4 Hz	0.03 Hz

Table 1. Key differences between the SRU and the visible CCD channel in MICAS.

A major challenge with the use of MICAS in place of the SRU is that in an arbitrary attitude, the probability of a detectable star being in the MICAS FOV is too low. The solution is to constrain the spacecraft to attitudes that satisfy one of two criteria: either one and only one preselected star of sufficient magnitude ( $m_v < 6$ ) is in the MICAS FOV or the duration at the attitude is short enough that gyros can be used. In contrast to the primary mission, the remainder of the extended mission could accommodate such a requirement, with four classes of attitudes needed: HGA on Earth, IPS thrusting to the comet, trajectory correction maneuvers, and science data acquisition at Borrelly.

NASA and JPL approved the ambitious and very risky second phase of the recovery. Because the primary mission had already concluded successfully, the consequences of a failure were deemed low. The likelihood of success, as perceived by the DS1 project and communicated clearly and frequently to managers at JPL and NASA, was also low. It was believed that with about two more months the probability of success would be significantly higher, but the opportunity to have a chance for a comet encounter led the project to pursue the more ambitious plan.

Work on a new system began in February. During four months, software and operational methods were designed, developed, tested, and integrated. In addition to testbed testing, some developmental tests were conducted on the spacecraft.

Although scattered light in MICAS had been investigated in detail during the primary mission, there were no data on the signature at some attitudes that would be important during operations with the new system. MICAS is aligned with the +z axis, and scattered light is independent of roll angle around the Sun-spacecraft line, so further measurements were possible even without three-axis knowledge or control. To verify that the scattered light models were correct in attitudes which had not previously been explored, the spacecraft was commanded to turn in the same way as before an HGA pointing session. Thus, in these tests the desired turn was accomplished by having ACS move the Sun by the desired angle from the center of the SSA. Images were collected and returned during HGA tracks. These data validated the scattered light models and yielded improved confidence in selecting appropriate parameter values for the new system.

Although MICAS was used frequently during the primary mission, both for validation of it as a technology and by AutoNav for its optical navigation images, it would be used much more extensively for the remainder of the extended mission. It includes no moving parts to wear out (part of the technology innovation), but to assure that it would be capable of providing reliable attitude information, in March a test began that acquired one MICAS image every 30 seconds and transferred it to the spacecraft computer. A problem which had been suspected from limited evidence in the testbed and in flight manifested itself in this long-duration test. On very rare occasions, several data words are dropped somewhere along the way from MICAS to the board that provides the interface to the spacecraft computer. This renders the image file useless and sometimes causes the transfer of subsequent images to stop. The problem happens so rarely that it had not shown up clearly before, but with images planned essentially continuously for the rest of the mission, it had to be accommodated. It was determined that the addition of a simple command before every image request would clear the problem; although the previous image would be lost if the words were missing, subsequent images could transfer normally.

The new system that was developed requires one and only one preselected bright star in the MICAS FOV. Initialization and

acquisition are discussed below, but once ACS has acquired that star, it tracks it by issuing a request to MICAS to take an image which is delivered to AutoNav for processing. AutoNav subtracts a background image to suppress some of the scattered light effects and locates all the candidate stars, some of which may be cosmic rays. Building upon the existing capability in AutoNav to process MICAS images was critical for timely completion of the software.

The locations and integrated intensities of candidate stars are delivered to ACS. Ground commands instruct ACS what the stellar magnitude is for the star to be tracked, and ACS identifies the star from among the candidates found by AutoNav. ACS includes limits on how much the observed magnitude is allowed to differ from the expected magnitude. This helps account for several effects, including certain regions in the MICAS FOV with greatly reduced optical throughput. Further discrimination is provided by using estimates of spacecraft motion (as measured by gyros and the SSA) to predict where the star should be in each image, based on where it was observed in earlier images. ACS then incorporates the measured location of the star into its control loop. The system has protections built in to accommodate a missed picture or a picture in which the star fails to show up.

To acquire a star, ACS is instructed to turn (relying on gyros and the SSA during the turn) to the target attitude. When it arrives at the estimated location, it begins a mosaic with MICAS. The mosaic size can be adjusted;  $3 \text{ FOV} \times 3 \text{ FOV}$  (with some overlap from each element to the next) is normally used. First it acquires the image to be used for background subtraction. Then at each element of the mosaic it collects two images and searches for the candidate star in each (to avoid confusion from cosmic rays). If a star close enough to the desired magnitude is found, the mosaicking is terminated and ACS transitions from acquisition to tracking. If no star is observed within that range, the mosaic continues until it is completed. At completion, if a star within a broader range was observed somewhere in the mosaic, it is used. If that criterion is not satisfied, a new mosaic is begun. The new mosaic can overlap the previous one or be moved to a new location by a desired angle, depending upon the values in

parameters which are easily updated by ground command.

The coordinates of the target star are included with its magnitude (and MICAS integration time) in the ground-generated commands; ACS assumes when it has found a star that is consistent with the observed magnitude and SSA angle that the star is the correct one.

Including a star catalog on board was considered and rejected during development. The very ambitious schedule led to a decision to limit the complexity of the on-board system and instead rely on ground tools to generate the commands to include all necessary information for each turn and subsequent acquisition.

It had been planned that new software would be loaded for the comet encounters during the extended mission. As a result, the DSN schedule already had adequate coverage in April 2000 for such a tracking-intensive activity. To allow more time for development (at the expense of less time for in-flight testing), the DSN agreed to postpone the tracking allocation. Coverage for uplinking the software began on May 30.

DS1 completely reloaded the flight software three times during the primary mission, in all cases to enable new technology validations. Thus the process for replacing the software was well understood. The principal differences between this new software load and previous ones were that uplink rates were lower (because of greater geocentric range -- the spacecraft was 2.0 AU from Earth during this load) and each pass that did not have a hand-over from a preceding pass had to begin with the time consuming process of bringing the HGA to Earth-point. To load the 4 megabytes of software required 267 command files.

On June 8, 2000 the computer was commanded to reboot and install the new software. By this time, the updated trajectory analysis, based on new operational principles discussed below, showed that to maintain adequate margin for reaching Borrelly, thrusting should resume by about July 5. As a result, an intensive test and verification campaign was necessary as soon as the new software was running on board. But because of the fast pace

of the work leading up to loading the software, in-flight tests had not been designed in detail. A rapid cycle of design, development, testbed testing, and spacecraft execution of flight software and operational tests was undertaken.

The new capabilities were activated and tested methodically but quickly. One feature of the new ACS is the use of all available data (from the SSA and, when tracking, MICAS images) to estimate gyro drifts. Thus, during the HGA passes even before locking to a star, the attitude stability, as revealed by the carrier power detected at the DSN, was much improved.

The first attempt to lock to a star was on June 12. The initialization discussed here represents a combination of what was executed then with the general procedure that was developed for subsequent use in the event that the spacecraft loses its attitude knowledge. The HGA is brought to Earth-point, but now with the gyro biases estimated. The star pattern near the MICAS FOV is predicted for the HGA being pointed to Earth and the Sun being at the known offset angle in the SSA. The dominant uncertainty in this is in the ability to determine the angle between the HGA boresight and the Earth-spacecraft line from the measured signal strength at the DSN. The accuracy with which this knowledge could be fed back to the spacecraft is estimated to be  $4^\circ$ ; one contribution to this error is the unpredictable component of the gyro drift during the round-trip light time.

A bright star or grouping of bright stars near the MICAS FOV is selected, and the spacecraft is commanded to turn to it and begin mosaicking. Once ACS finds a star and begins tracking it, telemetry shows the measured stellar magnitude. If that is insufficient to determine unambiguously which star is being tracked, a deep image is taken in place and downlinked. Such an image reveals fainter stars than the on-board system can detect and aids in making a positive star identification. If any ambiguity remains, ACS is commanded to stop tracking the star, a short turn is executed to a chosen location, an image is collected, and the spacecraft returns to resume tracking the star. This image shows nearby stars to confirm the attitude.

While most of this process is taking place, the spacecraft can remain locked to the star and thus is very stable. It can remain on the star as long as necessary. Once the star identification is complete, if the spacecraft is not locked to the planned star, the quaternion that corresponds to the actual star is uplinked. The spacecraft then has a complete and accurate knowledge of its three-axis attitude. Experience has shown that once this is complete, the reliability of locking to other preselected stars after turns, even when turning  $> 50^\circ$ , is extremely good.

Tests conducted included turns to new stars, methods to acquire the attitude in the event it is lost, and tuning of parameters to make the system more robust. On June 21, after a hiatus of more than 7 months, the IPS was turned on to test ACS' ability to achieve thrust vector control. All tests were completed with excellent results.

## RESUMPTION OF ROUTINE OPERATIONS

An important aspect of the recovery effort was the design of a new trajectory to reach Comet Borrelly. Some of the challenges of low-thrust mission design are described by Rayman et al.<sup>5</sup> The new trajectory however had to satisfy a new constraint: all thrusting had to be in attitudes with a preselected bright star in the MICAS FOV. Studies showed that the trajectories that assumed continuously variable thrust directions (implemented during the primary mission with AutoNav updating the thrust attitude every 12 hours) could be broken into a small number of discrete segments, each with thrusting in a fixed inertial direction. The 8 months of thrusting needed to reach Comet Borrelly could be accomplished with as few as 3 segments, although to make the design more flexible, 6 segments are used. Each segment uses one reference star, designated a "thruster", for ACS to track. Targeting control is achieved by adjusting the transition time from one thruster to the next.

The project had set for itself a very aggressive, success-oriented schedule that included the resumption of thrusting on July 5. That would allow reasonable margin for later unexpected losses of thrust. The testing in June went so well, however, that thrusting to the comet began on June 28.



A typical week of thrusting with the new system is very similar to thrusting before the SRU failure, although AutoNav is not used for this. There was not enough time to make the major changes in AutoNav that would have been required for it to implement the new operational procedures. To begin, the spacecraft turns to the thruster, and ACS acquires and tracks it. The IPS is activated by stored sequence and thrusts throughout most of the week. Throttle levels during the week are chosen in advance and commanded from a sequence. Shortly before the next scheduled DSN pass, the spacecraft stops thrusting and turns to a star (designated an "Earth star") that has been selected for that date as one that when MICAS is pointed at it, the HGA is close to Earth-point.

Because of the high rate of hydrazine expenditure during the phase 1 HGA pointing activities and some of the tests early in phase 2, the margin for hydrazine for attitude control for the remainder of the mission is small. To reduce hydrazine consumption, when the HGA is on Earth-point, the IPS thrusts, thus allowing ACS to control two axes with xenon instead of hydrazine. When Earth-pointing is close to the desired thrust attitude, the IPS is operated at a high throttle level; otherwise, it is operated at a very low throttle level. This scheme has significantly reduced hydrazine use. As it turns out, the thrust attitude is close to the Earth-point attitude for most of the period of deterministic thrusting.

The new system has proven to be very successful. On only one occasion between the first lock (on June 12) and August 31 did the spacecraft lose track of a star without relocating it on its own. Apparently because of very high solar activity, which affected a number of spacecraft, on July 16 there were too many false star candidates in its pictures, and ACS lost track of the star. It continued IPS thrusting in the desired attitude however, using the gyros and SSA. As it slowly drifted with the gyros, it mosaicked until it found a star that satisfied the criteria it had been applying. That stopped the drift, and allowed the gyro drift estimate to be refined, although the star was not the desired one. It eventually lost that star and found another one. Nevertheless, its estimates of gyro drift were accurate enough that it had not moved far from the thruster. When the next DSN track was scheduled, on July 18, it turned from its

thrusting attitude, which was treated as being correctly locked on the thruster, to the expected location of the Earth star. It began with the wrong star, so it turned to an incorrect attitude and could not locate the Earth star. But it was close enough that the HGA was on Earth-point. The operations team quickly discovered the problem. Following the procedure described above, the spacecraft was commanded to return a deep image, and later was directed to a nearby star. The lost thrust time in this case was negligible.

The resumption of long-term thrusting has enabled DS1 to set the record for the longest operating time of any propulsion system in space. The previous record was held by NASA's Space Electric Rocket Test (SERT) II, launched in 1970 and operated in Earth orbit. SERT II accumulated 3879 hours of operation on one of its two experimental ion engines before the engine failed from an internal short. On August 31, 2000, DS1 had more than 5300 hours of operation.

### MISSION PLAN

Each day of thrusting now consumes about 0.1 kg of Xe and yields a  $\Delta v = 7$  m/s. To maintain margin, the trajectory that is being followed does not assume thrusting for the 28 days around solar conjunction, on November 11, 2000. Any thrusting that can be accomplished then will increase mission margins further, but flying a profile that relies on thrusting during that period is unnecessarily risky. Because telecommunications will be difficult or impossible for about 3 weeks around conjunction, any loss of thrust would not be correctable promptly.

In February 2001 new software will be loaded to provide the spacecraft with capabilities needed for the encounter with Comet Borrelly 7 months later. This will include some changes in autonomous encounter pointing (to find the nucleus in the presence of the optically confusing coma) that had been designed and tested prior to the SRU failure but were not included in the June 2000 software load. The newer software will incorporate modifications to allow the new ACS system to track the comet. During the software load period, all IPS thrusting will be on Earth-point and will contribute to reaching the comet. In-flight

encounter rehearsals will be conducted shortly after the software load and in June 2001.

The trajectory plan completes deterministic thrusting in March, but well over a month, and in some cases (depending upon the date) up to 4 months, of lost thrust can be accommodated.

Trajectory correction maneuvers (TCMs) will be executed as the spacecraft approaches the comet. For TCMs that last longer than the time allowed on gyros, the maneuvers will be decomposed into components whose vector sum achieves the required correction, while each component is aligned with a reference star.

Closest approach to Comet Borrelly will occur on September 23, 2001 at 1.36 AU from the Sun, within days of the comet's perihelion. The baseline plan is for the spacecraft to pass about 1500 km from the nucleus on the Sun-nucleus line at 17 km/s. (The spacecraft was not designed to encounter a comet, so minimizing the risk of dust impacts to the spacecraft is a key criterion in selecting such a large distance.) Two of the technologies tested during the primary mission were compact science instruments, each with a broad range of measurement capabilities integrated into one small package. Images and infrared spectra can be obtained with MICAS and the plasma experiment for planetary exploration (PEPE) can sample the dynamics and composition of the rich plasma environment.

The planned encounter of DS1 will be on Borrelly's fifteenth recorded apparition since its discovery in 1904. With a period of 6.9 years, the comet has been extensively studied by many investigators. This Jupiter-family comet is moderately active with well-defined coma and tails.

Detailed encounter plans have not yet been developed, but it is expected that the MICAS infrared imaging spectrometer will yield diagnostic data between 1.2  $\mu\text{m}$  and 2.8  $\mu\text{m}$  on the volatiles and other species exposed on the nucleus' surface. In addition, panchromatic images will be used to map the three dimensional form of the nucleus and near-nucleus jets as well as other discrete morphologies in the coma.

PEPE will measure the flux of cometary ions and electrons from 8 eV to 33 keV as a function of energy and angle, and ion composition from 1 to 140 amu/e. The results of these measurements will provide the velocity distributions and basic plasma parameters (density, velocity, and temperature) of ions and electrons plus the composition of the cometary ions.

## CONCLUSION

Many future missions which were previously unaffordable now may be undertaken, with the cost and risk of using new technologies being substantially reduced because of the successful results of DS1's testing. Although the primary mission's only requirements were to assess the payload of high-risk technologies, the extended mission offers an opportunity to go beyond those objectives and to return important science data at Comet Borrelly. The failure of the SRU early in the extended mission made it appear likely that any further accomplishments from DS1 were unlikely, but the ambitious and successful recovery have returned the spacecraft to smooth operations.

## ACKNOWLEDGMENTS

The members of the DS1 mission operations and anomaly recovery team are gratefully acknowledged for their meritorious and even heroic work, upon which this overview is based. The support of JPL and NASA HQ management during this very difficult recovery effort is also greatly appreciated.

Some of the research described in this paper was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

## REFERENCES

1. Rayman, Marc D. and David H. Lehman, "NASA's First New Millennium Deep-Space Technology Validation Flight," Second IAA International Conference on Low-Cost Planetary Missions, Laurel, MD, April 16-19, 1996, IAA-L-0502.

2. Rayman, Marc D., Philip Varghese, David H. Lehman, Leslie L. Livesay, "Results from the Deep Space 1 Technology Validation Mission," *Acta Astronautica* **47**, p. 475 (2000).

3. Rayman, Marc D. and David H. Lehman, "Deep Space One: NASA's First Deep-Space Technology Validation Mission," *Acta Astronautica* **41**, p. 289 (1997).

4. Soderblom, L. A., "DS1 Micas Full-Disk-Rotation Spectra Of Mars (1.2-2.8 Microns)," Submitted to 32nd Annual AAS Division for Planetary Sciences Meeting, October 23-27, 2000, Pasadena, CA.

5. Rayman, Marc D., Pamela A. Chadbourne, Jeffery S. Culwell, and Steven N. Williams, "Mission Design for Deep Space 1: A Low-Thrust Technology Validation Mission," Third IAA International Conference on Low-Cost Planetary Missions, Pasadena, CA, April 27 - May 1, 1998, IAA-L98-0502.